

Estimation of recruitment in catch-at-age models

Mark N. Maunder and Richard B. Deriso

Abstract: Management strategies must be designed to take into account the uncertainty inherent in fish populations and their assessments. Annual recruitment variation is an important component of uncertainty. Several methods that allow the estimation of annual recruitment in statistical catch-at-age models are described: (a) maximum likelihood estimation with no penalty on the annual recruitment residuals, (b) maximum likelihood estimation with a lognormal penalty on the annual recruitment residuals, (c) importance sampling to numerically approximate the marginal likelihood with a lognormal penalty on the annual recruitment residuals, and (d) full Bayesian integration using Markov Chain Monte Carlo with a lognormal prior on the annual recruitment residuals. Simulation analysis is used to test the performance of these methods. All four methods perform similarly at estimating quantities that are based on averaging or summing multiple estimates of annual recruitment; however the marginal likelihood method (c) and Bayesian integration (d) perform best at estimating annual recruitment and the standard deviation in annual recruitment residuals (σ_R) when catch-at-age data is missing for some years. The ability to estimate σ_R can be important for defining uncertainty when developing management strategies. The methods are applied to a New Zealand snapper (*Pagrus auratus*) stock and the estimate of σ_R is approximately 0.6.

Résumé : Les stratégies de gestion doivent être élaborées de manière à tenir compte de l'incertitude inhérente aux populations de poissons et à leurs évaluations. La variation annuelle du recrutement est une composante importante de cette incertitude. On trouvera ici la description de plusieurs méthodes qui permettent d'estimer le recrutement annuel à l'aide de modèles statistiques basés sur les captures en fonction de l'âge : (a) une estimation de vraisemblance maximale sans pénalité pour les résiduels du recrutement annuel, (b) une estimation de vraisemblance maximale avec une pénalité de type lognormal sur les résiduels du recrutement annuel, (c) un échantillonnage pondéré pour obtenir une approximation de la vraisemblance marginale avec une pénalité de type lognormal sur les résiduels du recrutement annuel et (d) une intégration bayésienne complète à l'aide de la méthode de Monte Carlo par chaîne de Markov avec un a priori de type lognormal sur les résiduels du recrutement annuel. Une analyse de simulation permet de vérifier la performance de ces méthodes. Les quatre méthodes fonctionnent de manière semblable dans l'estimation de valeurs basées sur la somme ou la moyenne d'estimations multiples du recrutement annuel; cependant, la méthode de vraisemblance marginale (c) et l'intégration bayésienne (d) sont les meilleures pour estimer le recrutement annuel et l'écart type des résiduels du recrutement annuel (σ_R), lorsque les données de captures en fonction de l'âge sont absentes pour certaines des années. La capacité d'estimer σ_R peut être importante pour définir l'incertitude dans l'élaboration de stratégies de gestion. Ces méthodes sont appliquées à un stock de dorades royales (*Pagrus auratus*) et l'estimation de σ_R est d'approximativement 0,6.

[Traduit par la Rédaction]

Introduction

Annual recruitment variation is an important component for the management of fish stocks. Management strategies must be designed to take into account the uncertainty inherent in fish populations and their assessments. Uncertainty can be divided into estimation uncertainty and future uncertainty. Estimation uncertainty encompasses all uncertainty about the model structure and the associated parameter estimates. Future uncertainty includes how processes may change in the future irrespective of how well we have estimated

them in the past. Recruitment variation, a form of future uncertainty, can be a significant component of the total uncertainty, especially if the recruiting year class comprises a large proportion of the total population. Populations with different levels of recruitment variability have different probabilities of extinction and recovery. One method to determine uncertainty from recruitment variation is to assume that future recruitment will occur with a similar distribution and time-series structure to historical recruitment. Therefore, it is important to have good estimates of historic recruitment variability.

Is it appropriate to put distributional constraints on annual recruitment in a fisheries stock assessment model? Some analysts would argue that we should not include any prior information in the analysis and let the data provide information about the variability in recruitment. The problem with this argument is that many applications only have sporadic catch-at-age data for the early time period and the estimates of annual recruitment for this period fluctuate more than would be reasonable for a typical fish stock. The variability

Received 7 November 2002. Accepted 31 August 2003.
Published on the NRC Research Press Web site at
<http://cjfas.nrc.ca> on 31 October 2003.
J17174

M.N. Maunder¹ and R.B. Deriso. Inter-American Tropical Tuna Commission, 8604 La Jolla Shores Drive, La Jolla, CA 92037-1508, U.S.A.

¹Corresponding author (e-mail: mmaunder@iattc.org).

seen in the annual recruitment estimates is exaggerated by the uncertainty in these estimates, and any inferences based on the point estimates of recruitment would be invalid (e.g., stock–recruitment relationships, environmental correlations with recruitment, stochastic forward projections). The inclusion of estimates of uncertainty for the recruitments is essential for these analyses, but frequently they are not included. The above argument also assumed that we have not learned anything from the many years of fisheries research, which is inconsistent with the scientific method (Hilborn and Liermann 1998). From previous studies, we know that recruitment is usually lognormally distributed and that the average standard deviation of the logarithm of the annual recruitment residuals over many fish species is around 0.6 (Beddington and Cooke 1983).

The method that is often used to constrain annual recruitment variation in stock assessment models (e.g., Maunder and Starr 2001) is similar to methods used for random effects models (Davidian and Giltinan 1995). In contrast to the argument presented above, the methods used for random effects models make prior assumptions about the distribution of parameters to allow information to be shared between data sets. The annual recruitment is comprised of two components, a fixed effect that is common to all years and a random effect that is different for each year. The fixed effect is usually the average recruitment or the recruitment expected from a stock–recruitment relationship. The random effect is the annual residual and is constrained by a distributional assumption. The distribution that is used to constrain the random effect is usually lognormal with the standard deviation of the logarithm of the annual recruitment residuals fixed (often at 0.6, e.g., Smith and Punt 1998). In traditional random effects models, the standard deviation is estimated as a parameter in the model rather than being fixed at a predetermined value (Davidian and Giltinan 1995). The estimation methods used in fisheries stock assessment are often based on maximizing the full likelihood function (e.g., Smith and Punt 1998; Maunder and Starr 2001). However, the estimation methods used in traditional random effects models usually integrate over the random effects to produce a marginal likelihood function; this marginal likelihood is maximized to estimate the model parameters (Davidian and Giltinan 1995). The marginal likelihood is often a better choice because as the number of years in the model increases, the number of parameters does not increase. Therefore, asymptotic approximations work better for the marginal likelihood (Gelman et al. 1995). This is particularly true for annual recruitment residuals in years for which there are no catch-at-age data. Unfortunately, annual recruitment is nonlinearly intertwined in the likelihood functions for catch-at-age, catch-per-unit-effort (CPUE), and trawl survey data, so that analytical solutions for the integral are not possible. Therefore, numerical integration or an approximation is required to calculate the marginal likelihood function.

Full Bayesian integration is an alternative to maximizing the marginal distribution. Bayesian analysis integrates the probability distribution across all model parameters as opposed to just across the random effects. However, in Bayesian analysis, prior distributions are required for all model parameters. Bayesian analysis has been used in several fisheries applications (e.g., McAllister et al. 1994; Maunder et al.

2000; Maunder and Starr 2001) and has recently been reviewed by Punt and Hilborn (1997) and McAllister and Kirkwood (1998).

We use simulation analysis to investigate several estimation methods that incorporate annual recruitment residuals in a simple catch-at-age model: (a) maximum likelihood estimation with no penalty on the annual recruitment residuals, (b) maximum likelihood estimation with a lognormal penalty on the annual recruitment residuals, (c) importance sampling to numerically approximate the marginal likelihood (integrated across the annual recruitment residuals) with a lognormal penalty on the annual recruitment residuals, and (d) full Bayesian integration with a lognormal prior on the annual recruitment residuals. We also investigate the effect of using an informative prior for the standard deviation of the logarithm of the annual recruitment residuals for methods *b–d* and the effect of recruitment being uniformly distributed rather than lognormally distributed as assumed. We apply the methods to estimate the variation in recruitment for the snapper (*Pagrus auratus*) stock on the west coast of the North Island of New Zealand.

Estimation methods

Annual recruitment, R_t , is defined as average recruitment, \bar{R} , multiplied by an annual residual, ϵ_y , that is exponentiated and bias corrected (e.g., Maunder and Starr 2001):

$$(1) \quad R_t = \bar{R} e^{\epsilon_y - 0.5\sigma_R^2}$$

The annual residual is modified by a lognormal bias correction factor, $-0.5\sigma_R^2$, to ensure that \bar{R} is equal to the mean recruitment (rather than the median), where σ_R is the standard deviation of the logarithm of the annual recruitment residuals.

The annual recruitment is assumed to come from a lognormal distribution so that $\epsilon_y \sim N(0, \sigma_R)$. A penalty (eq. 2) is added to the negative log-likelihood function to constrain the annual recruitment residuals (ignoring constants):

$$(2) \quad -\ln P(\epsilon | \sigma_R) = \sum_y \ln(\sigma_R) + \frac{\epsilon_y^2}{2\sigma_R^2}$$

The estimates using the full likelihood function minimize the combined negative log-likelihood (the terms penalized negative log-likelihood or negative log-posterior probability are also used because of the distributional assumption included for the recruitment residuals) including, in a typical statistical catch-at-age model, the catch-at-age data (C@A) and the CPUE or trawl survey data (eq. 3). This method assumes that the annual recruitment residuals are unknown parameters to be estimated.

$$(3) \quad -\ln L(\mathbf{y}|\boldsymbol{\theta}) = -\ln L(\text{CPUE}|\boldsymbol{\varphi}, \boldsymbol{\epsilon}) \\ - \ln L(\text{C@A}|\boldsymbol{\varphi}, \boldsymbol{\epsilon}) - \ln P(\boldsymbol{\epsilon}|\sigma_R)$$

where $\boldsymbol{\varphi}$ are the model parameters, including the process (σ_R) and observation error variances, excluding the annual recruitment residuals ($\boldsymbol{\epsilon}$); $\boldsymbol{\theta} = \{\boldsymbol{\varphi}, \boldsymbol{\epsilon}\}$; and \mathbf{y} is the full data set $\mathbf{y} = \{\text{CPUE}, \text{C@A}\}$.

Traditional random effects models assume that the random effects (annual recruitment residuals) are random variables

(nuisance parameters), and therefore, they are integrated out to form a marginal likelihood function, $P(\mathbf{y}|\boldsymbol{\varphi})$. To produce the marginal likelihood, the full likelihood is integrated over $\boldsymbol{\varepsilon}$ (eq. 4).

$$(4) \quad L(\mathbf{y}|\boldsymbol{\varphi}) = \int L(\text{CPUE}|\boldsymbol{\varphi}, \boldsymbol{\varepsilon}) L(\text{C@A}|\boldsymbol{\varphi}, \boldsymbol{\varepsilon}) P(\boldsymbol{\varepsilon}|\boldsymbol{\varphi}) d\boldsymbol{\varepsilon}$$

Because of the nonlinear nature of the problem, a numerical solution or an approximation is needed for the integral. We use importance sampling (IS) to numerically approximate the integral. This method integrates across the random effects and therefore cannot be used directly to estimate the annual recruitments or other derived parameters that are functions of the annual recruitments. These parameters are estimated by fixing the variance parameters at the values estimated using the marginal likelihood (eq. 4) while estimating $\boldsymbol{\varepsilon}$ and remaining parameters using the full likelihood (eq. 3).

The previously mentioned methods only use the historical information that recruitment is usually lognormally distributed. There is also historical information on the value of the standard deviation for this distribution, $P(\sigma_R)$. This information can be combined with the negative log-likelihood function (in this case, we use the posterior probability notation (eq. 5)):

$$(5) \quad -\ln P(\boldsymbol{\theta}|\mathbf{y}) = -\ln L(\text{CPUE}|\boldsymbol{\varphi}, \boldsymbol{\varepsilon}) \\ - \ln L(\text{C@A}|\boldsymbol{\varphi}, \boldsymbol{\varepsilon}) - \ln P(\boldsymbol{\varepsilon}|\sigma_R) - \ln P(\sigma_R)$$

The prior can be generated by using a parametric representation of the frequency distribution of σ_R from many fish stocks that have estimates of annual recruitment. For example, Smith and Punt (1998) suggest that $\mu_{\sigma_R} = 0.6$ based on Beddington and Cooke (1983). Assuming $\sigma_{\sigma_R} = 0.2$, σ_R is normally distributed, and ignoring constants,

$$(6) \quad -\ln P(\sigma_R) = 0.5 \frac{(\sigma_R - 0.6)^2}{0.2^2}$$

Prior information is usually associated with full Bayesian integration, which requires integration over all model parameters (see Punt and Hilborn (1997) for a review of Bayesian analysis in fisheries stock assessment). However, it is not necessary to do full Bayesian integration to make use of prior information. For example, the mode of the joint posterior distribution, which can be found using a function optimizer, can be used as estimates of the model parameters (e.g., Maunder and Starr 2001). Statistical inference should not be solely based on one's ideological view (e.g., Bayesian vs. Frequentist), but based on practical methods that perform well (precise and accurate) and are robust to assumptions. Therefore, we test the penalized likelihood, IS, and full Bayesian integration methods that include a prior for σ_R . We now describe the details of each of the estimation methods.

Likelihood (MLE)

In this method, the full negative log-likelihood, excluding the annual recruitment penalty term,

$$(7) \quad -\ln L(\mathbf{y}|\boldsymbol{\theta}) = -\ln L(\text{CPUE}|\boldsymbol{\theta}) - \ln L(\text{C@A}|\boldsymbol{\theta})$$

is minimized to estimate all parameters of the model. In the implementation of this method, \bar{R} is not estimated, as it would be a redundant parameter, and instead of estimating $\boldsymbol{\varepsilon}$, the R_y are estimated directly and

$$(8) \quad \bar{R} = \frac{1}{T} \sum_y R_y$$

$$(9) \quad \sigma_R = \left(\frac{\sum_y \left[\ln(R_y) - \frac{1}{T} \sum_y \ln(R_y) \right]^2}{T-1} \right)^{1/2}$$

Penalized likelihood (PL)

In this method, the full negative log-likelihood, including the annual recruitment penalty term (eq. 3), is minimized to estimate all parameters of the model, including the annual recruitment residuals, which are assumed to be unknown parameters to estimate. This is equivalent to finding the mode of the posterior distribution of a hierarchical Bayesian model with locally uniform priors on all model parameters except the annual recruitment residuals.

Importance sampling (IS)

The marginal likelihood function can be approximated with numerical integration using IS (Gelman et al. 1995, pp. 307–308). The importance function is a multivariate normal based on the estimates of the recruitment residuals, $\boldsymbol{\varepsilon}$, and the corresponding variance-covariance matrix \mathbf{V} when the variance parameters are fixed. The algorithm alternates between updating the importance function (step 2) and performing the numerical integration (steps 3–4).

- (1) Choose initial values for $\boldsymbol{\varphi}$.
- (2) Use an iterative function minimizer to estimate $\hat{\boldsymbol{\varepsilon}}$ (and $\hat{\mathbf{V}}$), \bar{R} , and q , while keeping σ_R and σ_{CPUE} fixed at their current estimates using the full negative log-likelihood function (eq. 3).
- (3) Draw n sets of parameters, indexed by i , from a joint normal distribution $N(\hat{\boldsymbol{\varepsilon}}, \hat{\mathbf{V}})$.
- (4) Estimate $\hat{\boldsymbol{\varphi}}$: (a) use an iterative function minimizer to estimate $\hat{\boldsymbol{\varphi}}$; (b) for each function evaluation of the function minimizer, calculate the objective function as

$$-\ln \left[\frac{1}{n} \sum_i \frac{L(\text{CPUE}|\boldsymbol{\varphi}, \boldsymbol{\varepsilon}^i) L(\text{C@A}|\boldsymbol{\varphi}, \boldsymbol{\varepsilon}^i) P(\boldsymbol{\varepsilon}^i)}{g(\boldsymbol{\varepsilon}^i)} \right],$$

where $g(\boldsymbol{\varepsilon}^i) = N(\boldsymbol{\varepsilon}^i; \hat{\boldsymbol{\varepsilon}}, \hat{\mathbf{V}})$ and $P(\boldsymbol{\varepsilon}^i) = N(\boldsymbol{\varepsilon}^i; 0, \hat{\sigma}_R \mathbf{I})$.

- (5) Repeat steps 2–4 several times.
- (6) Use an iterative function minimizer to re-estimate $\hat{\boldsymbol{\varepsilon}}$ and $\hat{\mathbf{V}}$ conditioned on the current estimates of σ_R and σ_{CPUE} using the full negative log-likelihood function (eq. 3).

The estimates of $\hat{\boldsymbol{\varepsilon}}$ and estimates of the derived parameters that are functions of $\hat{\boldsymbol{\varepsilon}}$ are taken from step 6. Estimates of the other parameters are taken from step 4 the last time that it is completed.

We found that 10 iterations of steps 2–4 were sufficient for the algorithm to converge and 1000 samples in step 3 re-

stricted the variation in the estimates between iterations after convergence to acceptable levels. It is possible that the population can become negative for some of the samples of the recruitment residuals from step 3. This causes computational errors and we used a smooth function causing the population size to be always slightly above zero when it would normally be predicted to go below zero. In a standard stock assessment model, the objective function would be highly penalized by the difference between the new biomass level and the predicted level that is negative to aid the optimization routine. However, in the IS method, one random sample from step 3 may cause the biomass to go negative, causing divergent behavior. Therefore, we did not include this penalty, and its exclusion did not appear to cause any significant bias in the results.

Full Bayesian integration (BI)

The Markov Chain Monte Carlo (MCMC) method is used to sample from the posterior distribution (Gelman et al. 1995). We use the MCMC method supplied with the AD Model Builder software (Otter Research Ltd. 2000), which starts the MCMC algorithm at the mode of the joint posterior distribution and uses multivariate normal jumping rules based on the estimated variance-covariance matrix. One hundred thousand MCMC steps were used and every 100th step was saved. Other methods, such as sample importance resampling (McAllister and Ianelli 1997), could also be used. Parameter estimates are calculated as the mean of the marginal posterior distribution (i.e., averaged across all the other model parameters). Only the prior for ϵ is assumed to be informative, the priors for the other parameters are chosen to be uninformative or diffuse (Table 1).

Penalized likelihood (PL – σ_R), importance sampling (IS – σ_R), and full Bayesian integration (BI – σ_R) with a prior on σ_R

These methods are the same as PL, IS, and BI except eq. 5 is used which has a prior on σ_R .

Simulation analysis

We use simulation analysis to test the performance of each of the methods described above. A simple age-structured model (Appendix A) was set up to simulate a population for $T = 20$ years, starting from an unexploited equilibrium population (with no variation in the initial conditions). The effort trajectory was set to give contrast in the biomass time series; it increased in the first 10 years, decreased for the next 5 years, and stayed low for the last 5 years. The model was used to generate catch, CPUE, and catch-at-age data. The simulated recruitment was generated as being comprised of T random annual residuals with standard deviation = 0.6 and mean = 0. The standard deviation of the observation error in the CPUE index, σ_{CPUE} , was set at 0.3, and the sample size of the catch-at-age data was set to 50. The same age-structured model was then fit to the data to estimate the model parameters. The parameters estimated in the model were average recruitment (\bar{R}), the standard deviation of recruitment residuals (σ_R), the catchability coefficient (q), the standard deviation of the fit to the CPUE data (σ_{CPUE}), and

Table 1. Prior distributions used for the model parameters in the Bayesian integration method.

Parameter	Prior
\bar{R}	$U(0, 2\bar{R}^{\text{true}})$
q	$\ln U(-\infty, \infty)$
σ_{CPUE}	$\ln U(-\infty, \infty)$
σ_R	$\ln U(-\infty, \infty)$
ϵ	$N(0, \sigma_R^2)$

the annual recruitment residuals (ϵ). The estimation scheme was initiated with the true values for q and σ_{CPUE} . \bar{R} was set to one and a half times the true value, $\epsilon = 0$, and $\sigma_R = 1.0$. The simulation analysis was repeated 500 times for two different scenarios: (1) using catch-at-age data for all years and (2) using catch-at-age data for the last 10 years only. We also generate simulated data using random recruitment that is uniform $(0, 2\bar{R})$.

For each simulation, we report the estimates of \bar{R} , the biomass in the last year of the simulation as a ratio of the biomass in the first year of the simulation (B_T/B_1), R_y , σ_R , and σ_{CPUE} . These estimates are compared with the true values by calculating the relative error $\frac{\theta_i - \theta_i^{\text{true}}}{\theta_i^{\text{true}}}$ and reporting the 10,

50, and 90 percentiles and the average absolute error. We only tested methods PL – σ_R , IS – σ_R , and BI – σ_R using simulated data with no catch-at-age data for the first 10 years.

Application

The methods are applied to the snapper stock on the west coast of the North Island of New Zealand. The data consist of total catch from 1931 to 1998, a CPUE relative abundance index from 1974 to 1991, and catch-at-age data for 1975–1976, 1979, 1986–1987, and 1989–1998 (Nick Davies, National Institute of Water and Atmospheric Research, P.O. Box 1043, Whangarei, New Zealand, personal communication). We estimate the variation in annual recruitment.

Results

The relative error in the estimates of \bar{R} is similar for all the methods (Table 2). There is very little bias with or without the first 10 years of catch-at-age data, except for BI that gives a 8% positive bias when the first 10 years of catch-at-age data are not used. The average error (average of the absolute value of the relative error) is increased when the first 10 years of catch-at-age data are not used and is highest for the likelihood method.

The results for the relative error in the estimates of B_1/B_T are also similar for all methods (Table 3). There is very little bias with or without the first 10 years of catch-at-age data, except for PL, which gives a 5% positive bias when the first 10 years of catch-at-age data are not used, and BI, which gives a 4% and 6% positive bias with and without the first 10 years of catch-at-age data, respectively. Unlike the results for \bar{R} , the average error is not greatly increased when the

Table 2. Relative error in the estimate of average recruitment \bar{R} from four estimation methods.

	Catch-at-age data used for all years				Catch-at-age data not used for first 10 years			
	Median	10%	90%	Average absolute	Median	10%	90%	Average absolute
MLE	0.00	-0.07	0.08	0.05	0.02	-0.15	0.36	0.16
PL	-0.01	-0.06	0.06	0.04	0.02	-0.10	0.20	0.09
IS	0.00	-0.06	0.08	0.05	0.01	-0.12	0.19	0.10
BI	0.02	-0.04	0.08	0.04	0.08	-0.06	0.27	0.12

Note: MLE, maximum likelihood estimates using the full likelihood function; PL, penalized likelihood; IS, importance sampling; BI, Bayesian integration.

Table 3. Relative error in the estimate of biomass in the last year as a ratio of the biomass in the first year B_1/B_T from the four estimation methods.

	Catch-at-age data used for all years				Catch-at-age data not used for first 10 years			
	Median	10%	90%	Average absolute	Median	10%	90%	Average absolute
MLE	-0.02	-0.25	0.27	0.17	0.01	-0.20	0.28	0.15
PL	0.01	-0.16	0.22	0.12	0.05	-0.16	0.35	0.17
IS	0.01	-0.16	0.23	0.12	0.01	-0.17	0.23	0.13
BI	0.04	-0.13	0.26	0.13	0.06	-0.12	0.29	0.14

Note: MLE, maximum likelihood estimates using the full likelihood function; PL, penalized likelihood; IS, importance sampling; BI, Bayesian integration.

first 10 years of catch-at-age data are not used, except for the PL method.

The estimates of σ_R differ substantially between the different estimation methods (Table 4). When all catch-at-age data were used, the MLE method gave moderately positively biased estimates of σ_R , and when the first 10 years of catch-at-age data were not used, the estimates of σ_R became extremely positively biased. This bias was due to many of the recruitment estimates in the first few years, in which there are no catch-at-age data, being estimated as either very large or close to zero. The PL method is highly negatively biased and this bias increases when the first 10 years of catch-at-age data are not used. For some of the simulated data sets, σ_R was estimated to be zero, implying that the recruitment is the same for each year. The percentage of data sets in which σ_R was estimated to be zero was much higher when the first 10 years of catch-at-age data are not used (41%) compared with when all catch-at-age data were used (2%). The IS and BI methods performed much better at estimating σ_R when the first 10 years of catch-at-age data were not used compared with the MLE and PL methods. These methods also perform better when all catch-at-age data were used.

The estimates of σ_{CPUE} all have similar levels of average error for the different methods, but different biases (Table 5). The MLE, PL, and IS methods have a negative bias and BI has a smaller, but positive, bias.

Estimates of the relative error in the annual recruitment are similar for all four methods when the catch-at-age data were used for all years (Fig. 1). The variance in relative error increased for the last 5 years because these cohorts are only observed in a limited number of catch-at-age data sets. The MLE method shows a very small negative bias for all years, and the other methods show a small positive bias. When the first 10 years of catch-at-age data were not used,

these methods performed differently, particularly in the first half of the time period (Fig. 2). The MLE method cannot adequately estimate recruitment in the first few years. As mentioned previously, these recruitments are often estimated to be either unrealistically high or close to zero. The PL method is better than the likelihood method at estimating recruitment in the first few years but has higher average error in the recruitment estimates for the latter years. The IS method performs better than the MLE and PL methods. The IS method shows less average error for the first few years compared with the MLE method and less average error in the latter years compared with the PL method. The BI method performs similarly to the PL method.

When the true recruitment comes from a uniform distribution, the errors in the estimates of average recruitment, depletion, and σ_{CPUE} are all very similar to those from when the true recruitment comes from a lognormal distribution. There is higher average error in the recruitment estimates for all but the likelihood method (Fig. 3). For the scenario in which there is no catch-at-age data for the first 10 years, there is much more error in the estimates of recruitment for the PL, IS, and Bayesian methods compared with when the true recruitment is from a lognormal distribution (Fig. 4). The increase in error occurs mostly in the first 10 years.

The inclusion of a prior on σ_R did not change the ability to estimate \bar{R} or σ_{CPUE} . Estimates of depletion were only slightly improved (the average of the absolute value of the relative error is 0.16, 0.12, and 0.13 for the PL, IS, and BI methods, respectively). However, it did considerably decrease the average error in the estimates of σ_R (the average of the absolute value of the relative error is 0.56, 0.13, and 0.13 for the PL, IS, and BI methods, respectively). It also slightly reduced the error in the estimates of recruitment for the BI method (Fig. 5). A misspecified prior ($\mu_{\sigma_R} = 0.3$) gave simi-

Table 4. Relative error in the estimate of the standard deviation of the annual recruitment residuals from the four estimation methods.

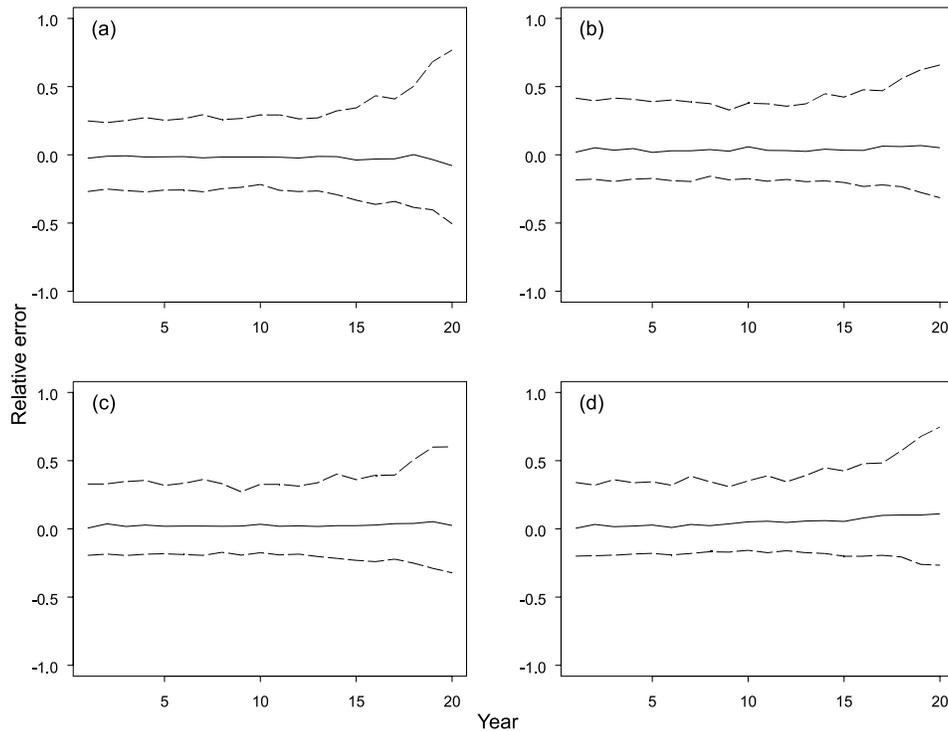
	Catch-at-age data used for all years				Catch-at-age data not used for first 10 years			
	Median	10%	90%	Average absolute	Median	10%	90%	Average absolute
MLE	0.09	-0.11	0.34	0.17	8.96	6.28	11.04	8.23
PL	-0.19	-0.45	0.04	0.23	-0.54	-1.00	-0.20	0.64
IS	-0.02	-0.22	0.21	0.14	-0.09	-0.35	0.18	0.19
BI	0.00	-0.23	0.23	0.15	-0.03	-0.33	0.28	0.20

Note: MLE, maximum likelihood estimates using the full likelihood function; PL, penalized likelihood; IS, importance sampling; BI, Bayesian integration.

Table 5. Relative error in the estimate of σ_{CPUE} from the four estimation methods.

	Catch-at-age data used for all years				Catch-at-age data not used for first 10 years			
	Median	10%	90%	Average absolute	Median	10%	90%	Average absolute
MLE	-0.05	-0.25	0.20	0.14	-0.11	-0.29	0.14	0.15
PL	-0.07	-0.26	0.16	0.14	-0.04	-0.24	-0.21	0.14
IS	-0.05	-0.24	0.19	0.14	-0.05	-0.24	0.19	0.14
BI	0.02	-0.19	0.27	0.14	0.02	-0.18	0.28	0.14

Fig. 1. Median (solid line), 10th percentile (lower broken line), and 90th percentile (upper broken line) for the relative error in annual recruitment for the (a) maximum likelihood estimate, (b) penalized likelihood, (c) importance sampling, and (d) Bayesian integration methods when using all the catch-at-age data.



lar results, causing a considerable negative bias in σ_R for all methods and a considerable negative bias in annual recruitment estimates for the BI method.

For the application to the snapper stock on the west coast of the North Island of New Zealand, the MLE method had problems converging. The PL method produced an estimate of zero. The IS method produced an estimate around 0.6.

The algorithm converged after about 10 iterations (Fig. 6). However, there is still variation in the estimates at each iteration. Increasing the number of samples from the importance function from 100 to 500 did not reduce this variation (Fig. 6). Therefore, a better estimate would be to take the average over a series of iterations rather than using the result after the final iteration. We use two averages, iterations 20

Fig. 2. Median (solid line), 10th percentile (lower broken line), and 90th percentile (upper broken line) for the relative error in annual recruitment for the (a) maximum likelihood estimate, (b) penalized likelihood, (c) importance sampling, and (d) Bayesian integration methods when not using the catch-at-age data for the first 10 years.

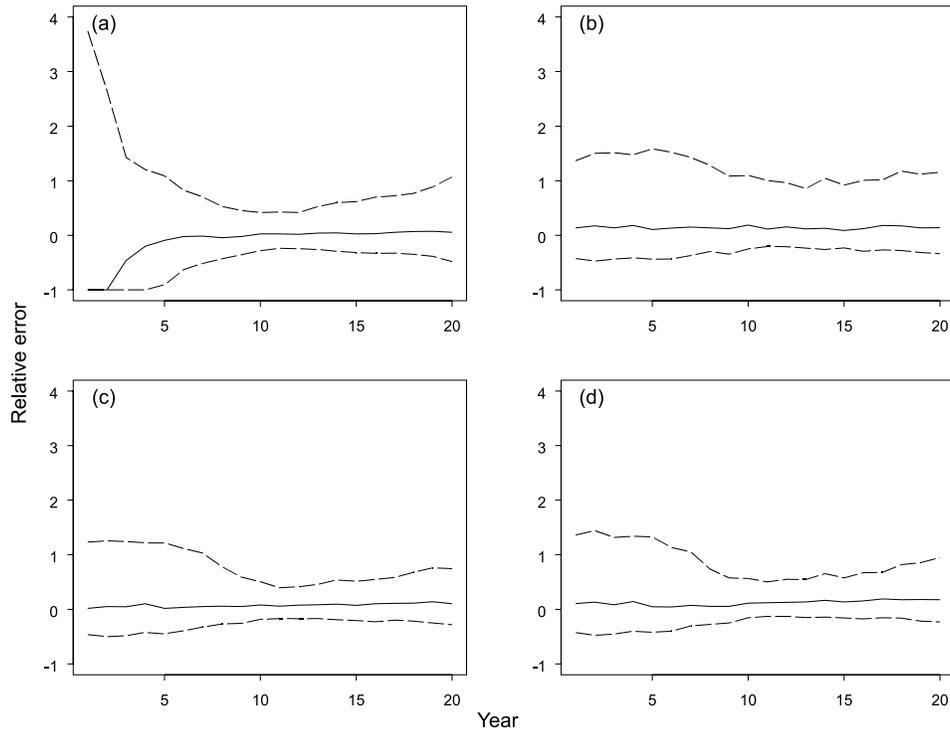


Fig. 3. Median (solid line), 10th percentile (lower broken line), and 90th percentile (upper broken line) for the relative error in annual recruitment for the (a) maximum likelihood estimate, (b) penalized likelihood, (c) importance sampling, and (d) Bayesian integration methods when using all the catch-at-age data and recruitment comes from a uniform distribution.

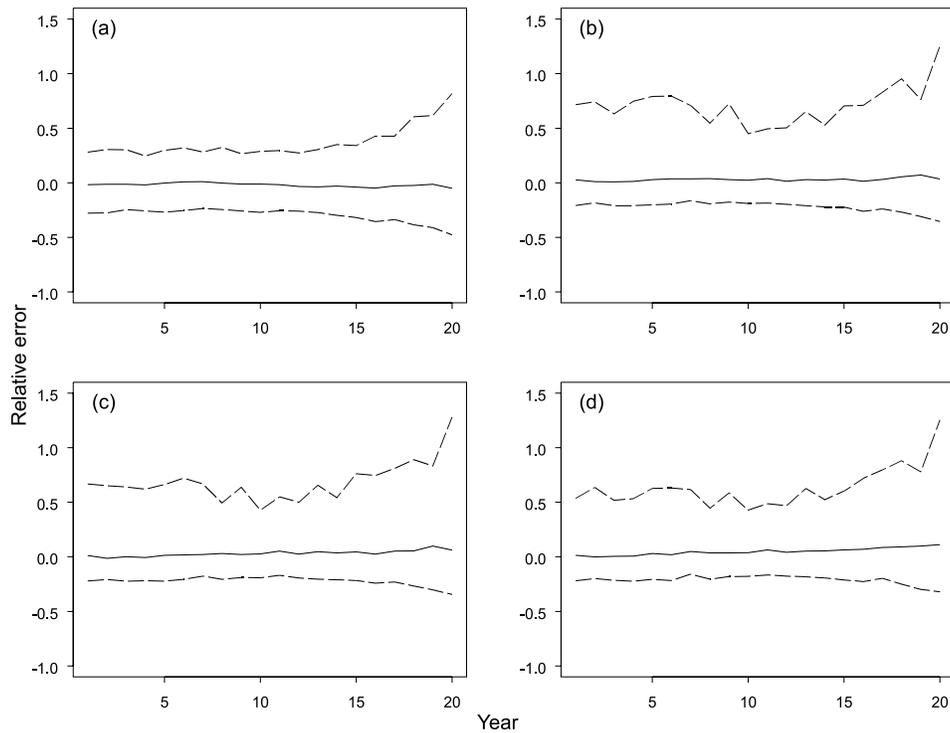
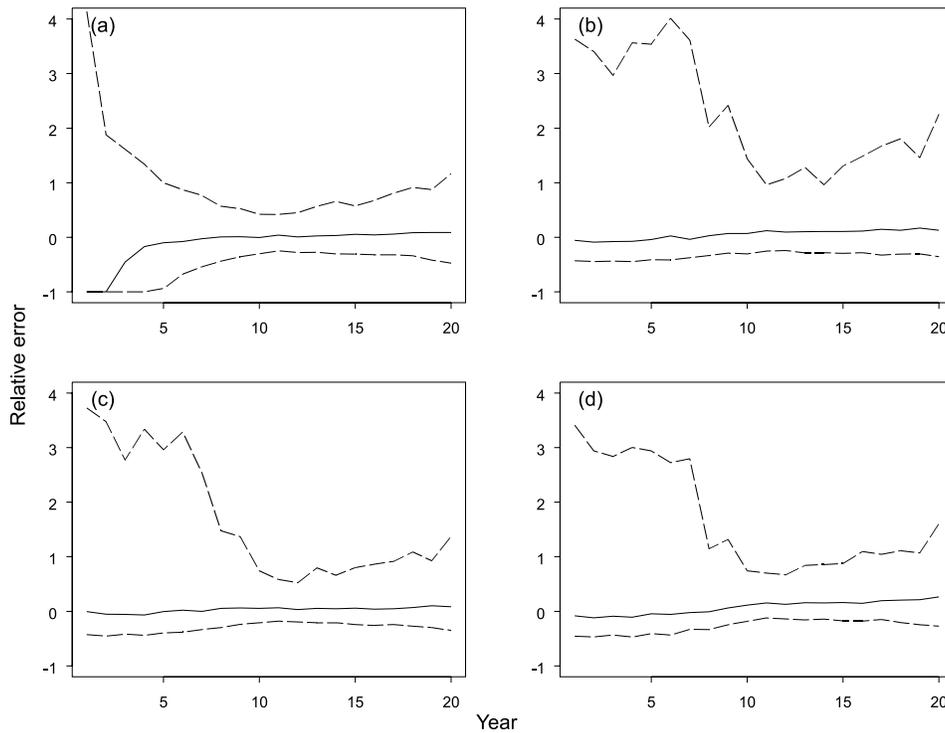


Fig. 4. Median (solid line), 10th percentile (lower broken line), and 90th percentile (upper broken line) for the relative error in annual recruitment for the (a) maximum likelihood estimate, (b) penalized likelihood, (c) importance sampling, and (d) Bayesian integration methods when not using the catch-at-age data for the first 10 years and recruitment comes from a uniform distribution.



to 100 and 80 to 100. For 100 samples from the importance function, the estimates of σ_R are 0.597 and 0.584, respectively. For 500 samples from the importance function, the estimates of σ_R are 0.620 and 0.584, respectively. The BI method estimates σ_R equal to 0.668 and 0.699 using the mode of the marginal distribution and the average, respectively. The posterior distribution ranges from about 0.4 to 1.0 and is skewed to the right (Fig. 7).

Discussion

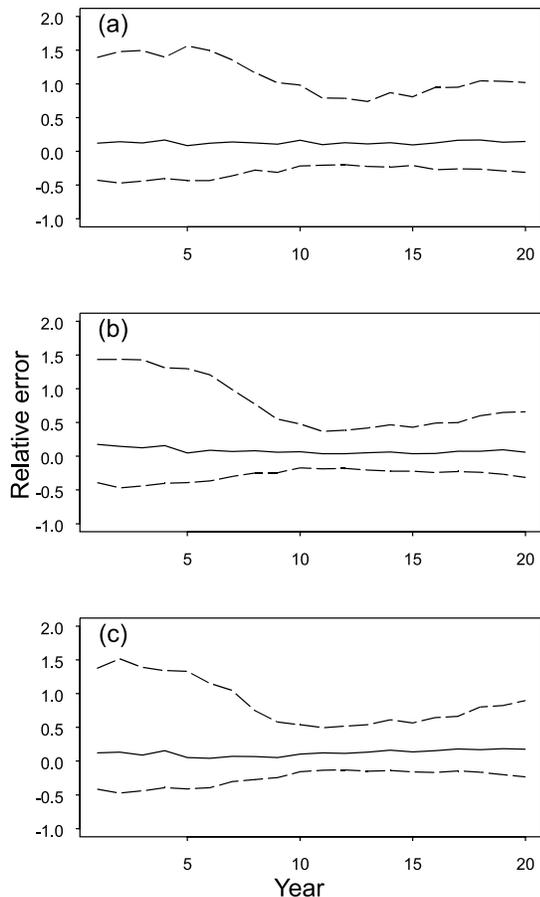
The methods presented in this paper for estimating annual recruitment in statistical catch-at-age models perform similarly when used to estimate quantities that are derived from a large number of annual recruitments (e.g., \bar{R}) or are a function of recruitments that have catch-at-age data associated with them (e.g., B_T/B_1). Average recruitment is unaffected by a lack of information (i.e., no catch-at-age data) for some annual recruitments because the CPUE time series covers the whole time period and the CPUE consists of the catches from several cohorts. Therefore, the individual annual recruitments can be incorrect, but the average of these recruitments must be close to the true average so that the predicted CPUE can be close to the observed CPUE. The methods also perform similarly for all quantities when the catch-at-age data are available for all years.

Annual recruitment and the standard deviation of the recruitment residuals are poorly estimated by the likelihood (MLE) and penalized likelihood (PL) methods when catch-at-age data are missing. The method based on the marginal likelihood (IS), which integrates across the recruitment re-

siduals, and the full Bayesian integration (BI) method perform much better in estimating the standard deviation of the recruitment residuals and annual recruitments. This indicates that any analyses that use the standard deviation of the recruitment residuals for forward projections or other analyses should use the marginal likelihood or full Bayesian integration, particularly if catch-at-age data are not available for the whole time period or if, for some years, the catch-at-age data are uninformative. For example, we suggest that if an environmental correlation with recruitment or a stock-recruitment model is being considered, it should be integrated into the stock assessment as suggested by Maunder and Starr (2001) and Maunder and Watters (2003) and a marginal likelihood or Bayesian integration used.

It is interesting to discover that σ_R could be estimated using the PL method for nearly all artificial data sets when catch-at-age data were available for all years. The estimation problem is more complex than a traditional random effects application, which has multiple measurements for a quantity of interest for each individual. Each annual recruitment residual has information about it contained in multiple catch-at-age data sets. In addition, because biomass is the accumulation of multiple year classes, each annual recruitment residual is represented in multiple CPUE data points. The fit to the catch-at-age data is based on the multinomial distribution, and therefore, the observation error variance for the catch-at-age data is determined by the sample size. The only observation error variance that is estimated is for the CPUE data. It is interesting to look at the negative log-likelihood profile for σ_R when using catch-at-age data for all years, which has a local minimum around the true value and a global mini-

Fig. 5. Median (solid line), 10th percentile (lower broken line), and 90th percentile (upper broken line) for the relative error in annual recruitment for the (a) penalized likelihood with a prior on annual recruitment residuals (σ_R), (b) importance sampling with a prior on σ_R , and (c) Bayesian integration with a prior on σ_R methods when not using the catch-at-age data for the first 10 years.



mum at zero (Fig. 8). The negative log-likelihood profile for σ_R when using catch-at-age only for the last 10 years only has a global minimum at zero (Fig. 8). However, if the estimates of $\sigma_R = 0$ are removed from the analysis, σ_R is still moderately negatively biased (median relative error is -0.17%) indicating that PL is still not a useful estimator for σ_R .

The results showed that estimates of annual recruitment have larger error if the true distribution of annual recruitment is different from that assumed in the analysis. However, the effect on the other quantities was small. The uniform distribution is an extreme sensitivity analysis compared with what empirical data suggest, therefore assuming a lognormal distribution is unlikely to cause large errors in the estimates for real applications.

Adding a prior on σ_R only improved the estimates of σ_R and not the other quantities. The PL method still estimated $\sigma_R = 0$ for a large proportion of the data sets. From these results, we conclude that putting a prior on σ_R is not particularly useful unless the data are very uninformative about annual recruitment. However, we did not investigate the influence on estimates of uncertainty (e.g., confidence intervals)

because the marginal likelihood method is too computationally intense.

The application to the snapper stock gave values of σ_R about 10–15% higher for BI compared with IS. The estimates from the BI method are sensitive to what measure is used to derive the estimates (e.g., average, mode of the marginal, median) and also probably to the prior distribution for the other model parameters. The estimates from the IS method are sensitive to which iteration is used as the estimate and the random seed. It is probably best to use an average over several iterations to get a more stable estimate. The estimates of σ_R are consistent with those suggested or assumed for other fish stocks.

The methods described in this paper have different computational demands and different estimation efficiencies. The MLE and PL methods only require the maximization (or minimization if the negative log-likelihood is used) to be carried out once. If an accurate optimizer is used (i.e., one based on automatic derivatives, e.g., Otter Research Ltd. 2000) and convergence has been achieved, this should be the exact (local) solution. The only complication is that a local rather than a global maximum may have been found. In contrast, the other methods are only approximations (to the marginal likelihood and the posterior distribution) and are more computationally demanding. The IS algorithm requires a numerical approximation based on averaging the likelihood from random samples of ϵ from the normal distribution $N(\hat{\epsilon}, \hat{V})$. Therefore, the likelihood function has to be evaluated for each sample. In addition, the IS method is also subject to the problem of local minima. The full Bayesian integration uses MCMC, which requires a large number of calculations of the objective function to approximate the posterior distribution. If too few iterations of the MCMC algorithm are performed, convergence may not be obtained, and if too many iterations are performed, the algorithm will take longer than needed. The priors used in the Bayesian analysis could be changed to see if other forms reduce the biases found for the Bayesian method when some of the catch-at-age data are missing. Medians or modes of the marginal posterior distributions could be used as alternative estimators. The MCMC algorithm started at the mode of the joint posterior distribution with jumping rules based on the covariance matrix. In initial simulations this caused bias in the estimates of σ_R because σ_R was estimated to be close to zero at the mode (see the PL results) and the estimated covariance matrix was inappropriate to cover the range of possible values for σ_R . To overcome this problem, we put a penalty on σ_R ($\mu_{\sigma_R} = 1.0$) while estimating the mode and the covariance matrix. The penalty was then removed when the MCMC algorithm was run. Sensitivity analyses suggested that this did not cause any significant biases in the results.

Simulated likelihood is a method that is becoming popular for implementing nonlinear random effects models (e.g., Harley 2002). Simulated likelihood involves taking random samples from the random effects distribution (this is the same as using the IS method with the importance sample equal to the random effects distribution). However, in applications like the one presented here in which the random effects parameters are combined in the model (i.e., the recruitments from multiple years are combined to determine the relative index of abundance and proportional catch-at-age data), all

Fig. 6. Estimates of σ_R (Rsd) for the snapper (*Pagrus auratus*) application after each iteration of the importance sampling method for two levels of the number of samples from the importance function (solid = 500, broken = 100).

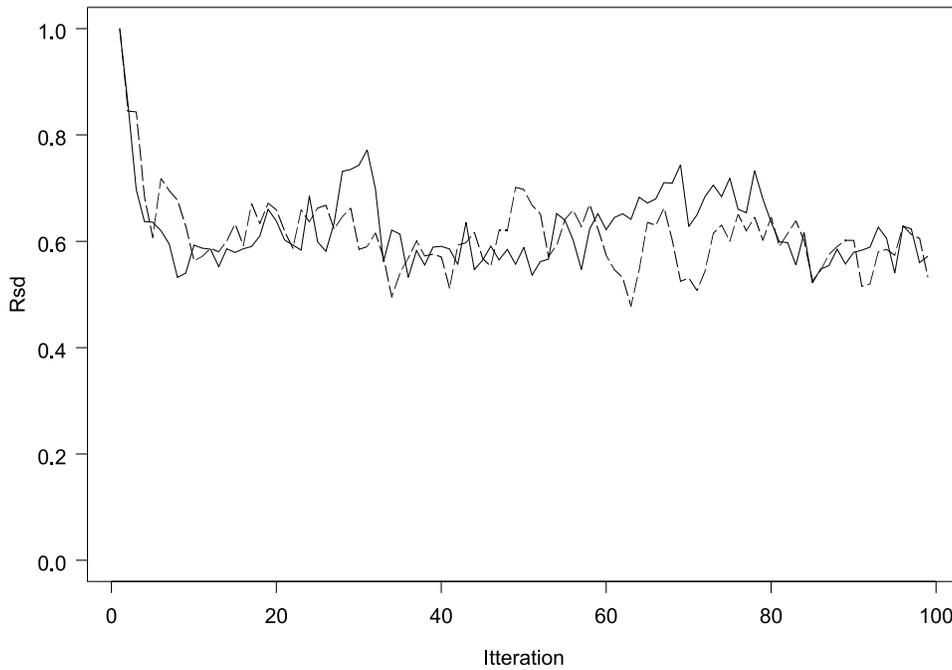
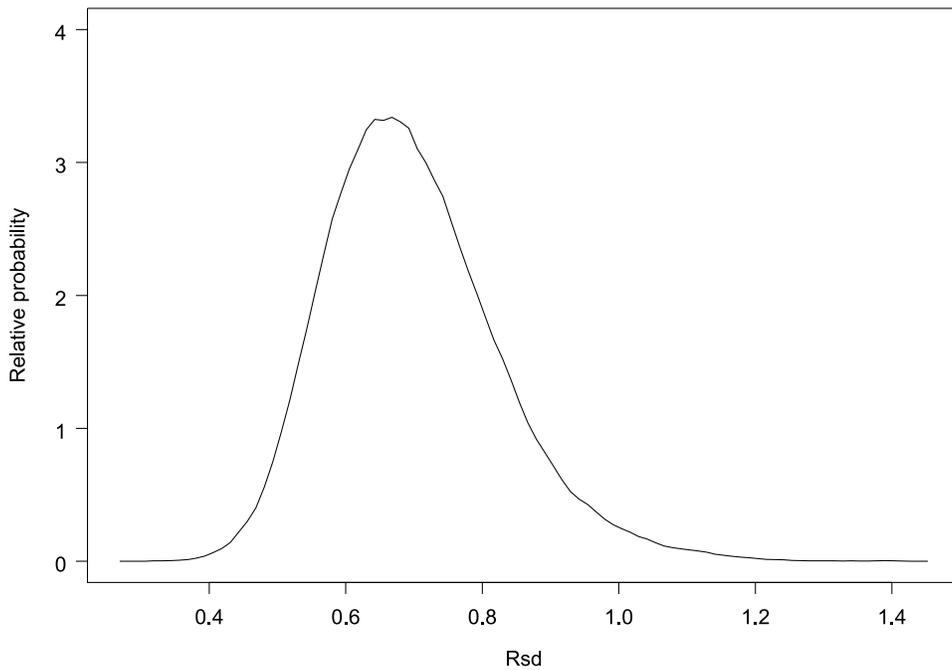


Fig. 7. Posterior distribution of σ_R (Rsd) for the snapper (*Pagrus auratus*) application from the Bayesian integration method.

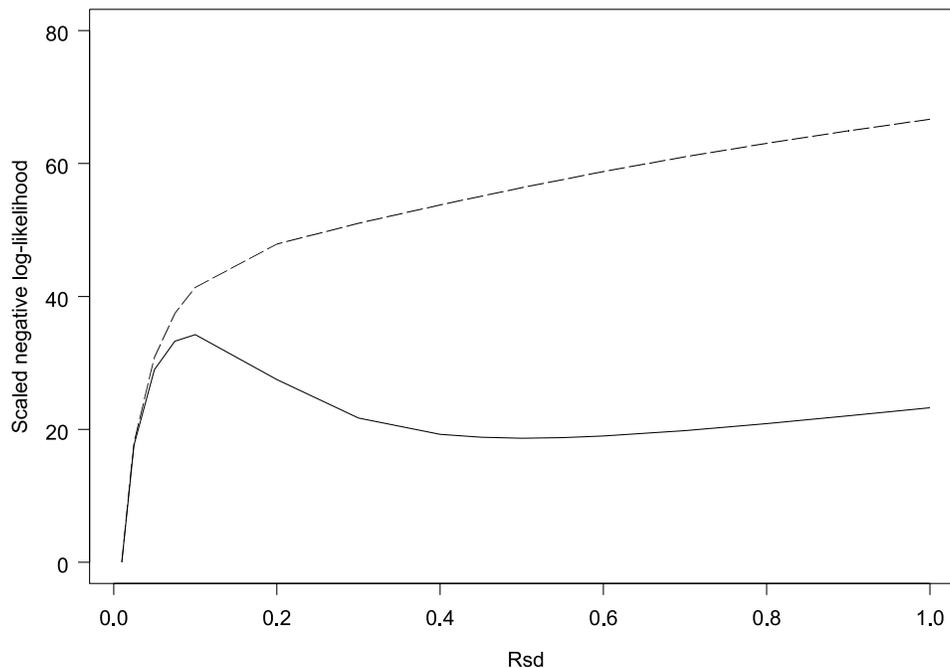


samples in an iteration have such small probability that the likelihood equals zero based on computer precision. The small probability occurs because the relative biomass in the last year, for a random sample of recruitments, is unlikely to be close to the observed index of abundance as it is a combination of all randomly sampled recruitments. Therefore, the minimization algorithm will not converge. This is our reason for using the IS algorithm. The IS algorithm generates samples that provide non-zero likelihoods because they are based

on the importance function that has been derived from fitting the model to the data and therefore the samples produce predictions close to the observed data. The IS method is similar to the SIR algorithm used in Bayesian analysis (see Punt and Hilborn 1997).

The estimation of σ_R is very important for projecting the uncertainty in the consequences of future management strategies. For example, the probability of extinction for a population with an average population growth rate of zero is

Fig. 8. Scaled negative log-likelihood for different values of σ_R (Rsd) from a single simulated data set with catch-at-age data for all years (solid line) and catch-at-age data only for the last 10 years (broken line).



higher for a population that has a high σ_R than for a population that has low σ_R . Therefore, it is important to differentiate between temporal variation in recruitment and estimation error. The results of our analyses show that using a marginal likelihood or full Bayesian integration should be the preferred methods when σ_R is used to determine the temporal variation in recruitment for forward projections. However, we have not investigated if the improved estimates of σ_R would produce better management advice. Given the greater computational demands of the marginal likelihood or full Bayesian integration, further studies should be carried out to determine if the additional computational demands are worth the benefit they provide. This would require carrying out management policy simulations.

The methods presented in this paper could be extended to include a stock–recruitment relationship, environmental relationship, or other model parameters. A stock–recruitment relationship or an environmental relationship can easily be included in the analysis by replacing \bar{R} in eq. 1 (see Maunder and Starr (2001) and Maunder and Watters (2003) for details). Process error for other model parameters can also be integrated out of the likelihood function. For example, fishing mortality is often modeled as a function of effort (E_y), catchability (q), and an annual residual (ε_y^E),

$$F_y = qE_y \exp(\varepsilon_y^E)$$

and a likelihood comparing the predicted and observed catch is included in the objective function. ε_y^E can be integrated out of the likelihood function in the same manner as ε_y is for recruitment. Meta-analysis, in which information is shared between stock assessments, could be carried out by modeling parameters (e.g., natural mortality) between populations as random effects and a marginal likelihood integrated across the random effects used for estimation.

In conclusion, if informative catch-at-age data are available for all years, then all of the methods perform well at estimating annual recruitment, average recruitment, and depletion level. In this case, the likelihood or penalized likelihood methods should be used because they are much less computationally intense, unless estimates of σ_R are needed. σ_R can be estimated using the penalized likelihood method if all years have informative catch-at-age data; however, σ_R is moderately negatively biased. The marginal likelihood and Bayesian integration methods perform best overall at estimating the model parameters when catch-at-age data are missing for some years.

Acknowledgements

Andre Punt provided comments on an earlier version of the manuscript and pointed out the local minima in the negative log-likelihood profile for σ_R .

References

- Beddington, J.R., and Cooke, J.G. 1983. The potential yield of fish stocks. *FAO Fish. Tech. Pap.* **242**: 1–47.
- Davidian, M., and Giltinan, D.M. 1995. *Nonlinear models for repeated measurement data*. Chapman and Hall, London.
- Fournier, D.A., Hampton, J., and Sibert, J.R. 1998. MULTIFAN-CL: a length-based, age-structured model for fisheries stock assessment, with application to South Pacific albacore, *Thunnus alalunga*. *Can. J. Fish. Aquat. Sci.* **55**: 2105–2116.
- Gelman, A., Carlin, J.B., Stern, H.S., and Rubin, D.B. 1995. *Bayesian data analysis*. Chapman and Hall, London.
- Harley, S.J. 2002. *Meta-analytical approaches to the study of fish population dynamics*. Ph.D. thesis, Dalhousie University, Halifax, N.S.

Hilborn, R., and Liermann, M. 1998. Standing of the shoulders of giants: learning from experience in fisheries. *Rev. Fish Biol. Fish.* **8**: 273–283.

Maunder, M.N., and Starr, P.J. 2001. Bayesian assessment of the SNA1 snapper (*Pagrus auratus*) stock on the northeast coast of New Zealand. *N.Z. J. Mar. Freshw. Res.* **35**: 87–110.

Maunder, M.N., and Watters, G.M. 2003. A general framework for integrating environmental time series into stock assessment models: model description, simulation testing, and examples. *Fish. Bull.* **101**: 89–99.

Maunder, M.N., Starr, P.J.S., and Hilborn, R. 2000. A Bayesian analysis to estimate loss in squid catch due to the implementation of a sea lion population management plan. *Mar. Mamm. Sci.* **16**: 413–426.

McAllister, M.K., and Ianelli, J.N. 1997. Bayesian stock assessment using catch-age data and the sampling – importance resampling algorithm. *Can. J. Fish. Aquat. Sci.* **54**: 284–300.

McAllister, M. K., and Kirkwood, G. P. 1998. Bayesian stock assessment and policy evaluation: a review and example application using the logistic model. *ICES J. Mar. Sci.* **55**: 1031–1060.

McAllister, M.K., Pikitch, E.K., Punt, A.E., and Hilborn, R. 1994. A Bayesian approach to stock assessment and harvest decisions using the sampling/importance resampling algorithm. *Can. J. Fish. Aquat. Sci.* **51**: 2673–2687.

Otter Research Ltd. 2000. An introduction to AD Model Builder version 4.5 for use in nonlinear modeling and statistics. Otter Research Ltd., Box 2040, Sidney, BC V8L 3S3, Canada.

Punt, A.E., and Hilborn, R. 1997. Fisheries stock assessment and decision analysis: the Bayesian approach. *Rev. Fish Biol. Fish.* **7**: 35–63.

Smith, A.D.M., and Punt, A.E. 1998. Stock assessment of gemfish (*Rexea solandri*) in eastern Australia using maximum likelihood and Bayesian methods. *In Fisheries stock assessment models. Edited by T.J. Quinn II, F. Funk, J. Heifetz, J.N. Ianelli, J.E. Powers, J.F. Schweigert, P.J. Sullivan, and C-I. Zhang. Alaska Sea Grant College Program, AK-SG-98-01. pp. 245–286.*

Appendix A. Description of simulator and estimator

The following is a description of the model equations used for the data simulator and for the estimator. The model is run from an unexploited state at the start of the fishery for 20 years. The model includes 10 age classes, with the 10th age class being a plus group.

Dynamics

$$N_{y,1} = \bar{R} \exp(\epsilon_y - 0.5\sigma_R^2)$$

$$N_{y,a} = (N_{y-1,a-1}(1 - u_{y-1}s_{a-1}))e^{-M} \text{ for } a < A$$

$$N_{y,A} = (N_{y-1,A-1}(1 - u_{y-1}s_{A-1}))e^{-M} + (N_{y-1,A}(1 - u_{y-1}s_A))e^{-M}$$

$$u_y = \frac{C_y^*}{B_y}$$

$$B_y = \sum_a N_{y,a} s_a w_a$$

where $N_{y,a}$ are the numbers in age class a at the beginning of year y , \bar{R} is the average recruitment, ϵ_y^R is the recruitment

anomaly for year y , σ_R is the standard deviation for the recruitment residuals, $C_{y,a}$ is the catch in numbers of age a individuals in year y , M is the natural mortality rate, A is the maximum age used in the analysis, u_y is the exploitation rate in year y , s_a is the selectivity to the fishing gear for age a individuals, C_y^* is the total catch in weight for year y , B_y is the exploitable biomass for year y , and w_a is the weight for an individual of age a .

Initial conditions

$$N_{1,a} = \bar{R}e^{-M(a-1)} \text{ for } 1 < a < A$$

$$N_{1,A} = \frac{N_{1,A-1}e^{-M}}{1 - e^{-M}}$$

Simulation

$$\epsilon_y \sim N(0, \sigma_R^2)$$

$$\epsilon_y^{CPUE} \sim N(0, \sigma_{CPUE}^2)$$

$$CPUE_y = qB_y \exp(\epsilon_y^{CPUE} - 0.5\sigma_{CPUE}^2)$$

$$n_{y,a} \sim \text{Multinomial} \left(\frac{C_{y,a}^{\text{numbers}}}{\sum_a C_{y,a}^{\text{numbers}}}, n = 50 \right)$$

$$C_{y,a}^{\text{numbers}} = N_{y,a} u_y s_a$$

where q is the catchability coefficient for the CPUE index, ϵ_a^{CPUE} is the observation error in the CPUE index, σ_{CPUE} is the standard deviation of the error in the CPUE index, $n_{y,a}$ is the number of individuals of age a in the catch-at-age sample in year y , and n is the number in the catch-at-age sample.

Estimation

Likelihoods

CPUE

$$-\ln L(\Pi|\theta) = \sum_y \left[\ln(\sigma_{CPUE}) + \frac{(\ln(CPUE_y) - \ln(qB_y))^2}{2\sigma_{CPUE}^2} \right]$$

Catch-at-age (modified from Fournier et al. 1998)

$$-\ln L(C@A|\theta) = -\ln \left[\exp \left\{ \frac{-(P_{y,a} - \hat{P}_{y,a})^2}{2\sigma_{C@A}^2} \right\} + 0.001 \right]$$

$$\sigma_{C@A}^2 = \frac{\left[(1 - P_{y,a})P_{y,a} + \frac{1}{A} \right]}{\sum_a n_{y,a}}$$

$$P_{y,a} = \frac{n_{y,a}}{\sum_a n_{y,a}}$$

$$\hat{P}_{y,a} = \frac{C_{y,a}^{\text{numbers}}}{\sum_a C_{y,a}^{\text{numbers}}} = \frac{N_{y,a} S_a U_y}{\sum_a N_{y,a} S_a U_y} = \frac{N_{y,a} S_a}{\sum_a N_{y,a} S_a}$$

where $n_{y,a}$ is the numbers of age a in the catch-at-age sample for year y .

Recruitment penalty

$$-\ln P(\epsilon^R | \sigma_R) = \sum_y \left[\ln(\sigma_R) + \frac{(\epsilon_y)^2}{2\sigma_R^2} \right]$$

Fixed parameters

$$w_a = l_a^3$$

$$l_a = 1 - e^{-0.1a}$$

$$M = 0.2$$

$$s = 1 \text{ for all } a$$